

10.03.00

A

10/02/00
JC929 U.S. PTO

UTILITY PATENT APPLICATION TRANSMITTAL

(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
POU920000173US1Total Pages in this Submission
37 (excluding references)

TO THE ASSISTANT COMMISSIONER FOR PATENTS

Box Patent Application
Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

COMMUNICATIONS BETWEEN PARTITIONS WITHIN A LOGICALLY PARTITIONED COMPUTER

and invented by:

Donald W. Schmidt et al.

JC929 U.S. PTO
10/02/00

10/02/00

If a **CONTINUATION APPLICATION**, check appropriate box and supply the requisite information: Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Which is a:

 Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Which is a:

 Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. Filing fee as calculated and transmitted as described below

2. Specification having 18 pages and including the following:
 - a. Descriptive Title of the Invention
 - b. Cross References to Related Applications (*if applicable*)
 - c. Statement Regarding Federally-sponsored Research/Development (*if applicable*)
 - d. Reference to Microfiche Appendix (*if applicable*)
 - e. Background of the Invention
 - f. Brief Summary of the Invention
 - g. Brief Description of the Drawings (*if drawings filed*)
 - h. Detailed Description
 - i. Claim(s) as Classified Below
 - j. Abstract of the Disclosure

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
POU920000173US1

Total Pages in this Submission
37 (excluding references)

Application Elements (Continued)

3. Drawing(s) (when necessary as prescribed by 35 USC 113)
 - a. Formal Number of Sheets _____
 - b. Informal Number of Sheets **5**
4. Oath or Declaration
 - a. Newly executed (*original or copy*) Unexecuted
 - b. Copy from a prior application (37 CFR 1.63(d)) (*for continuation/divisional application only*)
 - c. With Power of Attorney Without Power of Attorney
 - d. **DELETION OF INVENTOR(S)**
Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. 1.63(d)(2) and 1.33(b).
5. Incorporation By Reference (*usable if Box 4b is checked*)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. Computer Program in Microfiche (*Appendix*)
7. Nucleotide and/or Amino Acid Sequence Submission (*if applicable, all must be included*)
 - a. Paper Copy
 - b. Computer Readable Copy (*identical to computer copy*)
 - c. Statement Verifying Identical Paper and Computer Readable Copy

Accompanying Application Parts

8. Assignment Papers (*cover sheet & document(s)*)
9. 37 CFR 3.73(B) Statement (*when there is an assignee*)
10. English Translation Document (*if applicable*)
11. Information Disclosure Statement/PTO-1449 Copies of IDS Citations
12. Preliminary Amendment
13. Acknowledgment postcard
14. Certificate of Mailing

First Class Express Mail (*Specify Label No.*): **EK954544637US**

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
POU920000173US1

Total Pages in this Submission
37 (excluding references)

Accompanying Application Parts (Continued)

15. Certified Copy of Priority Document(s) *(if foreign priority is claimed)*

16. Additional Enclosures *(please identify below):*

Fee Calculation and Transmittal

CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	24	- 20 =	4	x \$18.00	\$72.00
Indep. Claims	3	- 3 =	0	x \$78.00	\$0.00
Multiple Dependent Claims (check if applicable)					\$0.00
				BASIC FEE	\$690.00
OTHER FEE (specify purpose)					\$0.00
				TOTAL FILING FEE	\$762.00

A check in the amount of _____ to cover the filing fee is enclosed.

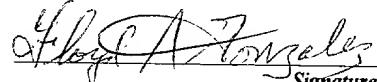
The Commissioner is hereby authorized to charge and credit Deposit Account No. **09-0463** as described below. A duplicate copy of this sheet is enclosed.

Charge the amount of **\$762.00** as filing fee.

Credit any overpayment.

Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.

Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).


Floyd A. Gonzalez
Signature

Floyd A. Gonzalez, Reg. 26,732
IBM Corporation
Intellectual Property Law
2455 South Road, P386
Poughkeepsie, New York 12601
(845) 433-1163
Fax: (845) 432-9786

Dated: **October 2, 2000**

cc:

Docket No:POU920000173US1

COMMUNICATIONS BETWEEN
PARTITIONS WITHIN A
LOGICALLY PARTITIONED
COMPUTER

APPLICATION FOR
UNITED STATES LETTERS PATENT

Express Mail Label No: EK954544637US

Date of Deposit: 0/2/2000

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee Service under 37 CFR 1.10 on the date indicated above and is addressed to Box Patent Application, Commissioner of Patents and Trademarks, Washington, D. C. 20231.

Susan L. Nelson 

I N T E R N A T I O N A L B U S I N E S S M A C H I N E S
C O R P O R A T I O N

COMMUNICATIONS BETWEEN PARTITIONS
WITHIN A LOGICALLY PARTITIONED COMPUTER

Field of the Invention:

5 The present invention relates to communications between processes in a multiprocessor system, and more particularly relates to communications between partitions within a logically partitioned computer.

Background of the Invention:

In a multiprocessor environment or a logically partitioned computer it is often desirable to move data from one processor to another, or from one partition to another. U.S. Patent 4,562,533 issued December 31, 1985 to Hodel et al. for DATA COMMUNICATIONS SYSTEM TO SYSTEM ADAPTER discloses a data processing system having multiple central systems and an adapter which is addressable by messages from each central system. The adapter provides handshake control, message synchronization and translation and the control of data being transmitted between a pair of central systems. The adapter includes a data buffer to provide intermediate storage of the information that is passed between the pair of central systems. A first message sequence is performed between a first, initiating central system and the adapter for transmitting information from the first central system to the adapter. A second message sequence is then performed between the adapter and a second, receiving central system for transmitting information from the adapter to the second central system.

20 Data may also be moved between servers in a physically partitioned computer. Data is typically moved to a shared intermediate area before it can be moved to the ultimate target server.

Summary of the Invention:

The present invention provides highly optimized any-to-any connectivity among discrete partitions or servers within a logically partitioned (LPAR) computer without requiring any physical cabling. Network latency is minimized because no physical I/O adapter is required to perform the desired data transfer among discrete servers within a computer. Instead, a direct 5 memory copy is performed by the sending central processing unit (CPU) from one server's memory to the memory of the other partition. Since the links among the discrete server are only virtual, no additional cabling or physical configuration is required when logical partitions are configured within the same computer. If this support is hidden under the TCP/IP protocol (as an internally implemented cluster Local Area Network (LAN)), then applications can gain 10 significant performance enhancements when communications occur via these virtual links, without any application changes. Security can be maintained among the partitions because the CPU I/O instruction is the only point of interception, since an adapter is not used for the communications. Since there is no physical media involved with these virtual links, the theoretical maximum bandwidth approaches that of the memory bus of the computer.

15 In the present invention, a centralized table defining the discrete servers within a computer is maintained in an area that is accessible by all the discrete servers. An I/O instruction detects that the target is connected via a virtual link. A lookup is performed within the centralized table to determine the target memory of the discrete server. Once the target is determined, the I/O instruction performs the data copy directly from the sender's buffer to the 20 previously queued target's receive buffer. The target partition may then be notified that the data is available in its receiver buffer by either a polling model, an interrupt model, or a combination of both.

25 It is an object of the present invention to provide a method and apparatus for transferring data from a discrete server in one logical partition to a discrete server in another logical partition while preserving security between the partitions.

It is another object of the present invention to provide a method and apparatus for establishing a send queue in a first logical partition, establishing a receive queue in a second

logical partition, registering the send queue in a lookup table available to all of the logical partitions, registering the receive queue in the lookup table, and executing a send instruction from the first partition which locates the send and receive queues in the lookup tables sends data in the send queue in the first logical partition to the receive queue in the second logical partition.

5 It is another object of the present invention to establish the send and receive queues in accordance with the direct input/output architecture.

It is another object of the present invention to store the lookup table in the hardware storage of the main storage of a data processing system.

10 It is another object of the invention to register the send queue using the logical partition identification and subchannel number of the first partition as the key.

It is another object of the invention to register the receive queue using the internet protocol address of the receive queue as the key.

15 It is another object of the invention to link the logical partition identification and subchannel number with internet protocol address with a queue control which points to the send queue and the receive queue.

It is another object of the invention to provide an indication in the hardware storage area which indicates when a send instruction from the subchannel of the first logical partition is to send data from the first partition to the second partition.

20 It is another object of the invention to provide a method an apparatus to include the IP address of the receive queue in the data stored in the send queue, execute a send instruction to search the lookup table for the logical partition identification and subchannel number to locate the send queue, to interrogate the data the send queue to find the IP address where the data is to

be sent, to search the lookup table for the IP address to find the receive queue, and to complete the data transfer from the located send queue to the located receive queue.

These and other objects will be apparent to one skilled in the art from the following drawings and detailed description of the invention.

5 Brief Description of the Drawings:

Fig. 1 is a schematic diagram of a network computing environment utilizing a channel subsystem usable with the present invention;

Fig. 2 is a schematic diagram of a single computer with shared physical memory and a plurality of discrete servers with a common lookup table of the present invention for transferring data from a sending discrete server to a target discrete server;

Fig. 3 is a schematic diagram illustrating the common lookup table of Fig. 2 including a hash tables control area, an source queue hash table, a target queue hash table, multiple queue controls, multiple QDIO queue sets, and means to add entries to the source queue hash table and target queue hash table;

Fig. 4 is a diagram of the hash tables control area of Fig. 3;

Fig. 5 is a diagram illustrating one of the queue controls of Fig. 3;

Fig. 6 is a diagram illustrating one of the queue sets of Fig. 3;

Fig. 7 is a diagram illustrating a send queue user buffer of the queue set of Fig. 6;

Fig. 8 is a diagram illustrating one of the entries of the source hash table of Fig. 3; and

Fig. 9 is a diagram illustrating one of the entries of the target hash tables of Fig. 3.

Description of the Preferred Embodiment:

An example of an existing data processing system architecture is depicted in Fig. 1. As shown in Fig. 1, information is passed between the main storage 110, and one or more input/output devices (hereinafter I/O devices) 190, using channel subsystems 150. Through the 5 switch 160, channel paths are established, comprising channels 155 and one or more control units shown at 180. These channel paths are the communication links established between the I/O devices 190 and the main storage for processing and exchange of information.

The main storage 110 stores data and programs which are input from I/O devices 190.

10 Main storage is directly addressable and provides for high speed processing of data by central processing units and one or more I/O devices. One example of a main storage is a customer's storage area and a hardware system area (HSA) to be discussed later. I/O devices 190 pass information to or from main storage via facilities provided in the channel subsystem 250. Some examples of I/O devices include card readers and punches, magnetic-tape units, direct-access storage devices (DASD), displays, keyboards, printers, teleprocessing devices, communication controllers and sensor-based equipment.

15 The main storage is coupled to the storage control element (SCE) 120 which in turn is coupled to one or more central processing units (CPU) 130. The central processing unit(s) is the control center of the data processing system and typically comprises sequencing and processing 20 facilities for instruction execution, initial program loading and other related functions. The CPU is usually coupled to the SCE via a bi-directional or unidirectional bus. The SCE, which controls the execution and queuing of requests made by the CPU and channel subsystem, is coupled to the main storage, CPUs and the channel subsystem via different busses.

25 The channel subsystem directs the flow of information between I/O devices and main storage and relieves the CPUs of the task of communicating directly with the I/O devices so that data processing operations directed by the CPU can proceed concurrently with I/O processing

operations. The channel subsystem uses one or more channel paths as the communication links in managing the flow of information to or from I/O devices. Each channel path consists of one or more channels, located within the channel subsystem, and one or more control units. In one preferred embodiment, a SAP I/O processor is also included as part of the channel subsystem.

5 As can be seen in Fig. 1, it is also possible to have one or more dynamic switches or even a switching fabric (network of switches) included as part of the path, coupled to the channel(s) and the control unit(s). Each control unit is further attached via a bus to one or more I/O device(s).

10 The subchannel is the means by which the channel subsystem provides information about associated I/O devices to the central processing units; the CPUs obtain this information by executing I/O instructions. The subchannel consists of internal storage that contains information in the form of a channel command word (CCW) address, channel path identifier, device number, count, status indications, and I/O interruption subclass code, as well as information on path availability and functions pending or being performed. I/O operations are initiated with devices 15 by executing I/O instructions that designate the subchannel associated with the device.

20 The execution of input/output operations is accomplished by the decoding and executing of CCWs by the channel subsystem and input/output devices. A chain of CCWs (input/output operations) is initiated when the channel transfers to the control unit the command specified by the first channel command word. During the execution of the specified chain of I/O operations, data and further commands are transferred between the channel(s) and the control unit(s).

25 Fig. 2 is a schematic diagram of a single computer with shared physical memory 210, and may be an IBM z/Series z/900 computer available from International Business Machines Corporation of Armonk, New York which is a follow-on computer of the IBM S/390 computer. The computer is divided up into a number logical partitions 212a -212n, each partition having discrete servers 214a-214n, respectively, labeled in Fig. 2 as discrete server 1 to discrete server n. Each discrete server has a TCP/IP layer 216a-216n, respectively, for handling the transmission

protocols for transmitting data in Input/Output (I/O) operations for networks, as is well known. Under each TCP/IP layer 216a-216n is a device driver 218a-218n, respectively, for driving data transmissions between the discrete servers, as will be discussed.

In the present invention, each device driver is similar to device drivers which drive the devices 190 of Fig. 1. However the device drivers 218 of Fig. 2, rather than driving I/O devices, drive data exchanges between the LPAR partitions, as will be explained. Each device driver 218 has a send queue 222, and a receive or target queue 220; the send queue 222 being used for sending data from the respective discrete server 214 when that discrete server is the sending server, and the receive queue 220 for receiving data for its respective discrete server 214 when that discrete server is the target server in a send operation, as will be described in connection with Fig. 3. A common lookup table 224 is in the HSA portion 225 of the main storage 110 of the single computer 210 across the entire computer, as explained in Fig. 1. This common lookup table 224 is a centralized table defining the discrete servers 214a-214n within the computer 210 and is maintained in HSA 225 that is accessible by all the discrete servers 214a-214n. However, the discrete servers can only register in the common lookup table using I/O type commands, and cannot retrieve any information from the lookup table 224, thus maintaining security between the servers.

Each device driver 218 is associated with a subchannel control block 227 which contains control information for the subchannel. As is known, the subchannel control blocks exist in HSA 225 and are uniquely identified by a subchannel number. The subchannel control block 227 includes an internal queued direct I/O (IQDIO) indicator 228 which indicates if this subchannel is an IQDIO subchannel. The IQDIO indicator 228 may be set by the channel path identifier (CHPID) definition statement during the configuration process, as is well known in the art.

The architecture of the computer 210 of the present invention adheres to the queued direct I/O (QDIO) architecture, as explained in U.S. Patent Application Serial No. 09/253,246 filed February 19, 1999 by Baskey et al. for A METHOD OF PROVIDING DIRECT DATA PROCESSING ACCESS USING A QUEUED DIRECT INPUT-OUTPUT DEVICE, owned by 5 the assignee of the present invention and incorporated herein by reference.

Fig. 3 is an illustration of the common lookup table 224 of Fig. 2, and includes hash tables control area 300, a source queue hash table 310, and a target queue hash table 320. The source queue hash table includes multiple entries starting with the first entry 311, each entry acting as a source queue duplicate list head (containing a pointer to duplicate list entries 312). 10 The target hash table 320 includes multiple entries starting with the first entry 321, each entry acting as a target queue duplicate list head (containing a pointer to duplicate list entries 322). A common queue control area 330 is shared by both send (using table 310) and receive (using table 320) processing. It will be noted that multiple 322s can point to a single 330. Each queue control 330 is linked to a QDIO queue set 340. New entries in the source queue hash table 310 are created at 312, and new entries in the target queue hash table 320 are created at 322, as will be explained.

Fig. 4 is a block diagram of the hash table control 300 and includes a hash table shared serialization lock 401, and a hash table exclusive update lock. Fig. 5 is a block diagram of the queue control 330 and includes a QDIO pointer 430 which points to the queue set 340, an 20 outbound lock 431, and an inbound lock 432.

Fig. 6 is a block diagram of the queue set 340 of Fig. 3 and includes a send queue 440 having multiple entries, and a receive queue 445 having multiple entries. The queue set 340 also includes a storage list status block (SLSB) 442 which shows the status of each entry in the send queue 440, and a storage list status block (SLSB) 447 which shows the status of each entry in the 25 receive queue 445. Each active entry of the send queue 440 has an associated buffer pointer 441 which points to a user buffer 443 for containing the data to be sent to the target LPAR partition. Fig. 7 is an illustration of the transfer data in the user buffer 243, and includes the target IP

address 244 to which the data is to be sent. Each active entry in the receive queue 445 is associated with a buffer pointer 446 which points to a user buffer 448 which is to receive the data transferred from the user buffer 443.

Fig. 8 is a block diagram illustrating the entries of the source queue hash table list 310 as set up at 312. Each entry includes the LPAR-ID.SUBCHANNEL# 410 used as a key to the table 311, the status 411 of the entry, the queue control pointer 412 which points to the control 330 for this entry, a next pointer 413 which points to the next entry 312 in the source hash table 310, and a previous pointer 414 which points to either the first entry 311 in the source hash table 310 or the previous entry created at 312. Similarly, Fig. 9 is a block diagram illustrating the entries of the target queue hash table as set up at 322. Each entry includes the IP address 420 used as a key to the table 321, the status 421 of the entry, a queue control pointer 422 which points to the control 330 for this entry, a next pointer 423 which points to the next entry 322 in the target hash table 320, and a previous pointer 424 which points to either the first entry 321 in the target hash table 320 or the previous entry created at 322.

The first step in transferring data from one LPAR partition to another, is to register a source or send queue 222 (represented in Fig. 2 as a downward arrow, and also shown as queue 440 in Fig. 6) and a receive or target queue 220 (represented in Fig. 2 as an upward arrow, and also shown as queue 445 in Fig. 6) for a send transaction. The registration process includes two steps: the first is to register the QDIO queue set 340 (one send queue 222 and one target queue 220) in the source queue hash table 310; and the second is to associate one or more IP addresses with the previously defined QDIO set 340 by adding entries to the target queue hash table 320. As each QDIO queue set 340 contains both a send queue 222 and a receive queue 220, both types of hash entries resolve into a single queue control structure 330 that contains a pointer to the QDIO defined queues

The source queue hash table registration is as follows:

- a. obtain the exclusive update lock 402 for the hash tables. Updates to both types of hash tables can be serialized with a single lock.

5

- b. using the LPAR-ID.SUBCHANNEL# as key into the source hash table 310, determine the appropriate duplicate list header location 311 in the source queue hash table 310.
- c. once found, use the pointers 413 and 414 in a well known fashion to scan all hash key 10 duplicate entries for an exact match with the LPAR-ID.SUBCHANNEL# being added. If found, then return the Duplicate Found error return to the TCP stack for the error to be dealt with there.
- d. if there are no duplicates, at 312, add an entry to the source queue hash table 310.
- e. create the queue control 330 that is to be associated with the newly created entry.
- f. release the exclusive update lock 402 for the hash tables.

The target queue hash table registration is as follows:

15

- a. obtain exclusive lock 402 for the hash tables. Again, updates to both types of hash tables can be serialized with a single lock.
- b. using the target IP address as the key, determine the appropriate duplicate list header location in the target queue hash table 321.
- c. once found. use the pointers 423 and 424 in a well known fashion to scan all hash key 20 duplicates for an exact match with the target IP addresses being added. If a duplicate is found, then return a Duplicate Found error to the TCP stack for the error to be handled there.
- d. if no duplicates are found, at 322, add an entry to the target queue hash table 321.
- e. using the LPAR-ID.SUBCHANNEL# from the input, perform a search of the source queue hash table 310 to find the previously defined queue control 330 that is to be associated with the newly created entry.
- f. release the exclusive update lock 402 for the hash tables.

25

A send operation to send data from one LPAR partition to another is as follows:

- a. As part of the processing of a socket API, the device driver 218 (software) modifies the send queue 440 (shown as downward arrow 222 in Fig. 2) to prime it with data to be transferred.

- b. a send is initiated by a SIGA instruction to the device driver 218. This SIGA instruction explained in the aforementioned 09/253,246 application includes the subchannel number associated with the send queue 222.
- c. the IQDIO indicator 228 of the subchannel control block 227 for the designated subchannel indicates that this is a IQDIO subchannel and that the send operation is to use the queue set 340 associated with this subchannel.
- d. the shared serialization lock 401 is obtained for the queue lookup table 224 access.
- e. the LPAR-ID from which the SIGA instruction is issued and the subchannel number in the instruction is used to build the LPAR-ID.SUBCHANNEL# key into the source hash table 310.
- f. obtain the outbound lock 431 to obtain exclusive serialization of the queue control 130 for the located entry in the source hash table 310.
- g. search the SLSB 442 to find the primed outbound storage buffer access list (SBAL) (shown as the buffer pointer 441) which points to the storage buffer access list element (SBALE) describing the packet of data to be moved to the target IP address.
- h. using the located SBAL, extract the destination IP address 244 from the outbound user buffer 443.
- i. use the IP address 244 to search the target queue hash table 320 to find the table entry 322 for the queue descriptor of the receive queue 220/445.
- j. obtain the inbound lock 432 to obtain exclusive serialization of the queue control 330 associated with the located target hash table entry 322.
- k. The SLSB 447 of the receive queue 445 is searched to find an empty SBAL to receive the data.
- l. move the data in user buffer 443 of the send queue 440 to the user buffer 448 of the receiver queue 445 using internal millicode mechanism that overrides the normal restrictions on data moves between storage addresses in different LPAR partitions.
- m. update the SLSB 442 of the send queue 440 and the SLSB 447 of the receive queue 445. These updates are visible to the software and allows program manipulation of the send and receive queues 222 and 220.
- n. release the shared serialization lock 401.

- o. set a program initiative of either a polling paradigm or a program interrupt, or some combination, for the partition that contains the receive queue 220 to indicate that new elements or data are available on the receive queue 220. Having been thus informed, software in the target partition may process the data in its receive queue 220.

5 It will be understood that in the present embodiment, steps b-o of the send operation are performed by hardware, making the performance of these steps very reliable and at hardware speed. However, these steps, or some portion of them, could be done in software, if desired. This invention may also be used to transfer data between multiple virtual servers within a single partition.

10 While the preferred embodiment of the invention has been illustrated and described herein, it is to be understood that the invention is not limited to the precise construction herein disclosed, and the right is reserved to all changes and modifications coming within the scope of the invention as defined in the appended claims.

Claims:

What is claimed is:

1. 1. In a data processing system having multiple logical partitions, a method of sending data from a first logical partition to a second logical partition comprising:
 2. establishing a send queue in the first logical partition;
 3. establishing a receive queue in the second logical partition;
 4. registering the send queue in a lookup table available to all of the logical partitions;
 5. registering the receive queue in said lookup table; and
 6. executing a send from the first logical partition which interrogates the lookup table to locate said send queue and said receive queue and sends the data in the send queue in the first logical partition to the receive queue in the second logical partition.
2. 2. The method of claim 1 wherein said send queue and said receive queue are established in accordance with queued direct input/output (QDIO) architecture.
3. 3. The method of claim 1 wherein said lookup table is stored in a hardware storage area of the data processing system.
4. 4. The method of claim 1 wherein input/output (I/O) operations for each logical partition is conducted via a subchannel associated with each logical partition, each logical partition having an identification, and the send queue is registered in said lookup table using the logical partition identification and subchannel number as a key.
5. 5. The method of claim 4 wherein said receive queue has an internet protocol (IP) address and said receive queue is registered in said lookup table using the IP address of the receive queue as the key.

1 6. The method of claim 5 further comprising linking the registration using said logical
2 partition identification and subchannel number with the registration of said IP address by a queue
3 control which points to said send queue and said receive queue.

1 7. The method of claim 6 further comprising setting an indication in the hardware storage
2 area indicating that the send instructions from the subchannel of the first logical partition is to
3 send data from said send queue of the first logical partition to the receive queue of the second
4 logical partition, and said send instruction interrogates said indication.

1 8. The method of claim 7 further comprising:

2 priming said send queue with data to be transferred, said data including the IP address of
3 the receive queue to which said data is to be transferred;

4 searching said lookup table using the logical partition identification and subchannel
5 number as a key to locate the send register registered in the lookup table with that key;

6 interrogating the data in said send queue to find the IP address of the receive queue;

7 searching said lookup table using said IP address as a key to locate the receive queue
8 registered in the lookup table with that key; and

9 sending the data in the located send queue to the located receive queue to complete the
10 transfer.

1 9. In a data processing system having multiple logical partitions, an apparatus for sending
2 data from a first logical partition to a second logical partition comprising:

3 means for establishing a send queue in the first logical partition;

4 means for establishing a receive queue in the second logical partition;

5 means for registering the send queue in a lookup table available to all of the logical
6 partitions;

7 means for registering the receive queue in said lookup table; and

8 means for executing a send from the first logical partition which interrogates the lookup
9 table to locate said send queue and said receive queue and sends the data in the send queue in the
10 first logical partition to the receive queue in the second logical partition.

1 10. The apparatus of claim 9 wherein said means for establishing said send queue and said
2 receive queue establish said send queue and said receive queue in accordance with queued direct
3 input/output (QDIO) architecture.

1 11. The apparatus of claim 9 further comprises a hardware storage area of said data processing
2 system for storing said lookup table.

1 12. The apparatus of claim 9 further comprising a subchannel associated with each logical
2 partition for conducting input/output (I/O) operations for the associated logical partition, each
3 logical partition having an identification, and said means for registering said send queue
4 registering said send queue in said lookup table using the logical partition identification and
5 subchannel number as a key.

1 13. The apparatus of claim 12 wherein said receive queue has an internet protocol (IP) address
2 and said means for registering said receive queue in said lookup table registering said receive
3 queue using the IP address of the receive queue as the key.

1 14. The apparatus of claim 13 further comprising queue control for linking the registration
2 using said logical partition identification and subchannel number with the registration of said IP
3 address, said queue control pointing to said send queue and said receive queue.

1 15. The apparatus of claim 14 further comprising an indication in the hardware storage area
2 for interrogation by said send instruction, said indication when set indicating when send
3 instructions from the subchannel of the first logical partition is to send data from said send queue
4 of the first logical partition to the receive queue of the second logical partition.

1 16. The apparatus of claim 15 further comprising:
2 means for priming said send queue with data to be transferred, said data including the IP
3 address of the receive queue to which said data is to be transferred;

4 means for searching said lookup table using the logical partition identification and
5 subchannel number as a key to locate the send register registered in the lookup table with that key;
6 means for interrogating the data in said send queue to find the IP address of the receive
7 queue;
8 means for searching said lookup table using said IP address as a key to locate the receive
9 queue registered in the lookup table with that key; and
10 means for sending the data in the located send queue to the located receive queue to
11 complete the transfer.

1 17. In a data processing system having multiple logical partitions, an apparatus for sending
2 data from a first logical partition to a second logical partition comprising:

3 a send queue in the first logical partition;

4 a receive queue in the second logical partition;

5 a lookup table available to all of the logical partitions, said lookup table having registered
6 therein said send queue and said receive queue; and

7 a device driver in said first logical partition which executes a send for interrogating the
8 lookup table to locate said send queue and said receive queue, and sending the data in the send
9 queue in the first logical partition to the receive queue in the second logical partition.

1 18. The apparatus of claim 17 wherein said send queue and said receive queue are established
2 in accordance with queued direct input/output (QDIO) architecture.

1 19. The apparatus of claim 17 further comprising a hardware storage area in said data
2 processing system for storing said lookup table.

1 20. The apparatus of claim 17 further comprising a subchannel associated with each logical
2 partition for conducting input/output (I/O) operations for the associated logical partition, each
3 logical partition having an identification, and said send queue is registered in said lookup table
4 using the logical partition identification and subchannel number as a key.

1 21. The apparatus of claim 20 wherein said receive queue has an internet protocol (IP) address
2 and said receive queue is registered in said lookup table using the IP address of the receive queue
3 as the key.

1 22. The apparatus of claim 21 further comprising queue control for linking the registration
2 using said logical partition identification and subchannel number with the registration of said IP
3 address, said queue control pointing to said send queue and said receive queue.

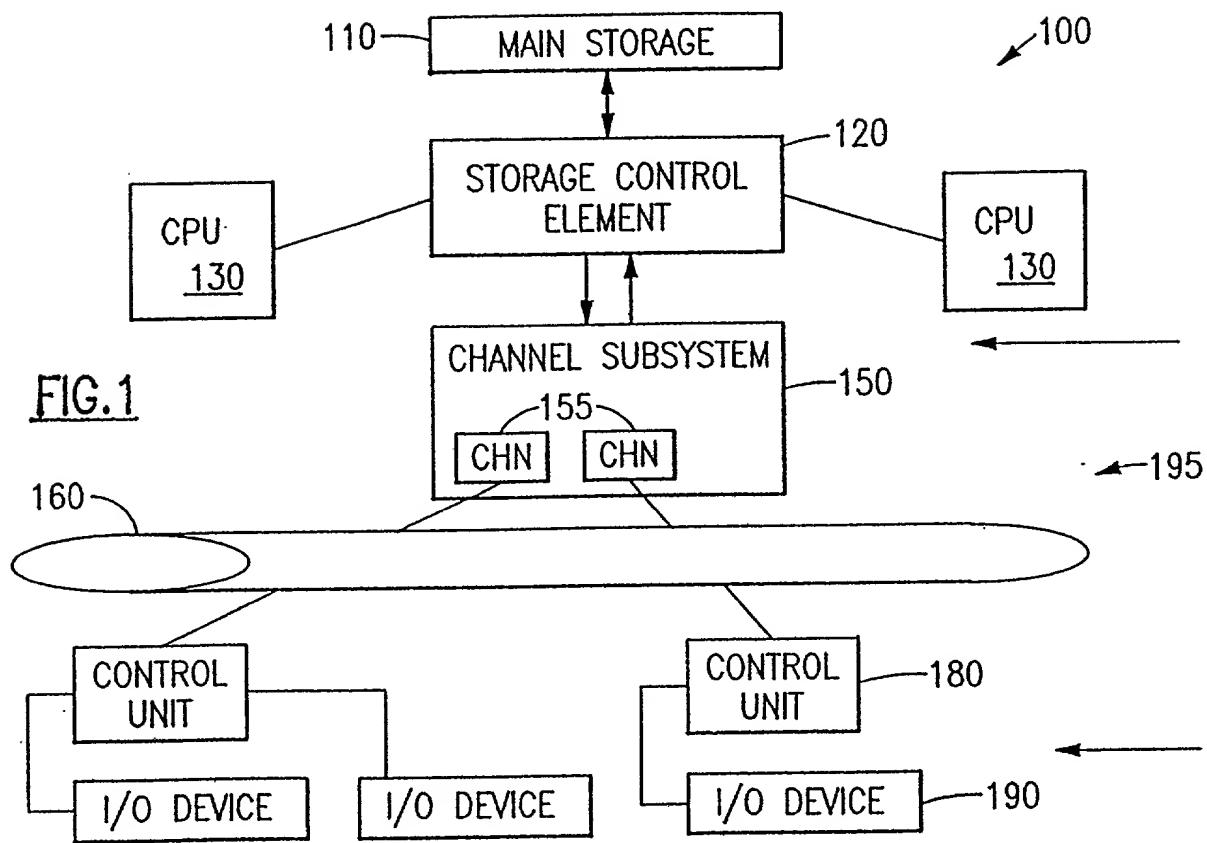
1 23. The apparatus of claim 21 further comprising an indication in the hardware storage area
2 for interrogation by said send instruction, said indication when set indicating when said send
3 instruction from the subchannel of the first logical partition is to send data from said send queue
4 of the first logical partition to the receive queue of the second logical partition.

1 24. The apparatus of claim 23 further comprising:
2 data to be transferred in said send queue, said data including the IP address of the receive
3 queue to which said data is to be transferred;
4 hardware searching said lookup table using the logical partition identification and
5 subchannel number as a key to locate the send register registered in the lookup table with that key;
6 hardware interrogating the data in said send queue to find the IP address of the receive
7 queue;
8 hardware searching said lookup table using said IP address as a key to locate the receive
9 queue registered in the lookup table with that key; and
10 means in said device driver for sending the data in the located send queue to the located
11 receive queue to complete the transfer.

COMMUNICATIONS BETWEEN PARTITIONS
WITHIN A LOGICALLY PARTITIONED COMPUTER

Abstract of the Disclosure:

Method and apparatus for sending data from one partition to a second partition within a logically partitioned computer. In a data processing system having multiple logical partitions, a send queue is established in the first logical partition, and a receive queue is established in the second logical partition. The send queue is registered in the send queue in a lookup table available to all of the logical partitions. The send queue is registered using as a key the logical partition identification of the first logical partition and the subchannel number (LPAR-ID.SUBCHANNEL#) of the subchannel assigned to the partition. The receive queue is registered in the lookup table using as a key, the internet protocol address of the receive queue in the second partition. A send instruction from the first logical partition is executed which interrogates the lookup table using the LPAR-ID.SUBCHANNEL# key to locate the send queue and IP address key to locate the receive queue, and sends the data in the send queue in the first logical partition to the receive queue in the second logical partition. This method and apparatus provides that discrete servers may be used in each logical partition, and data may be transferred between while maintaining security between the logical partitions.



POC9 20000173051
215

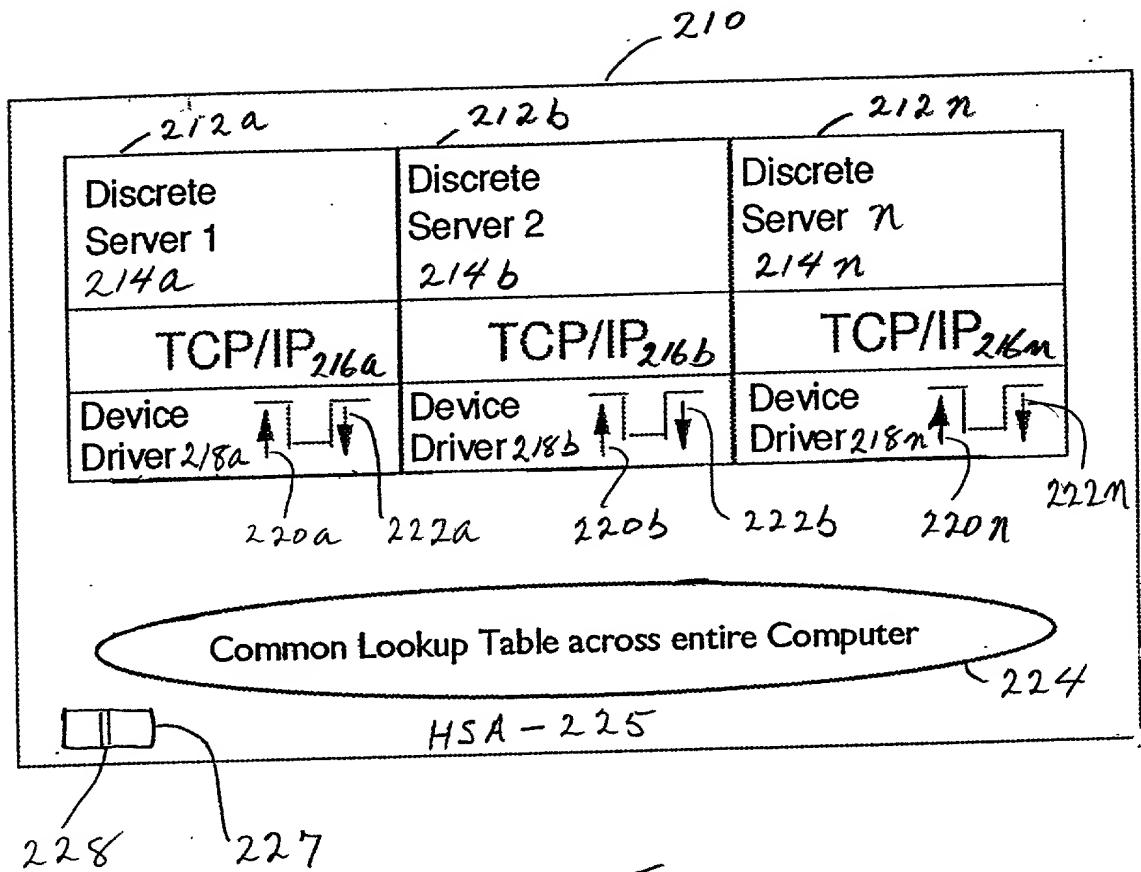
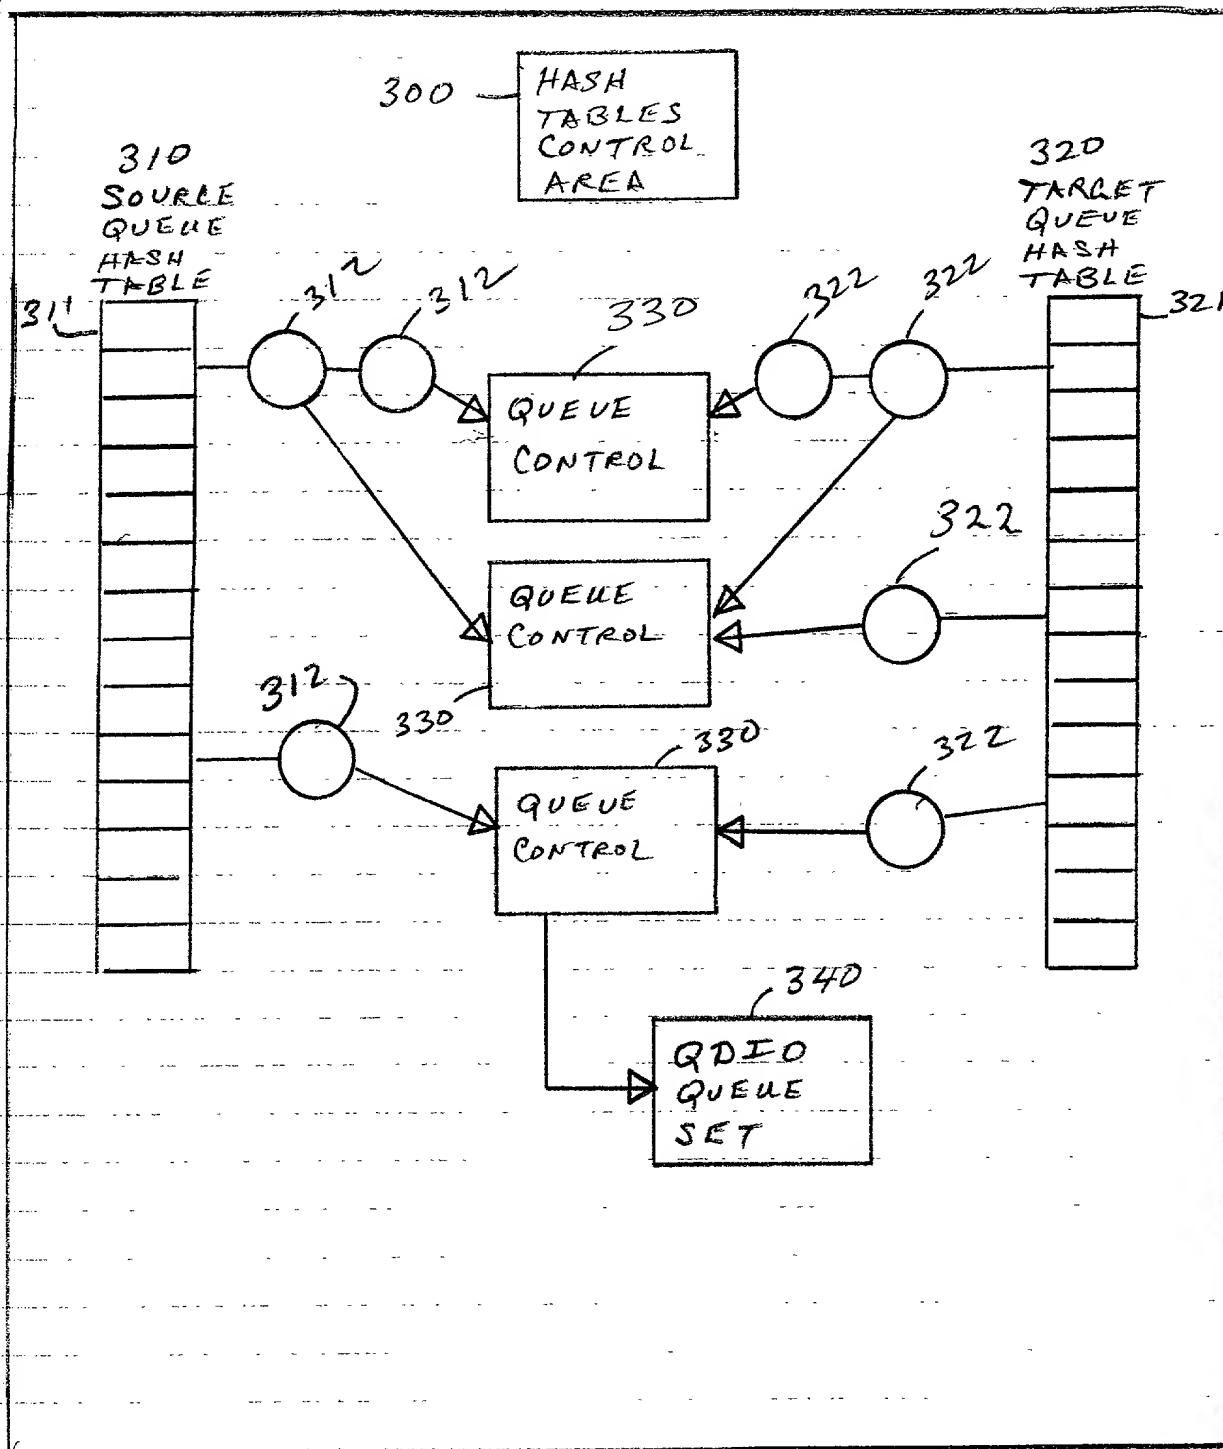


FIG 2

POU920000173051
3/5



224

FIG. 3

PC0920000173051
415

300 - HASH TABLE CONTROL

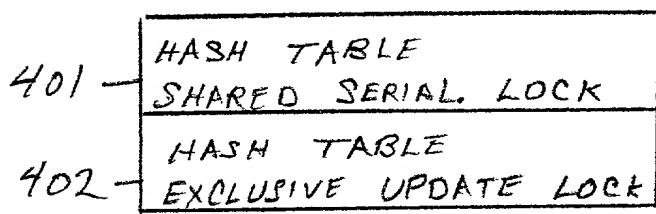


FIG. 4

330 - QUEUE CONTROL

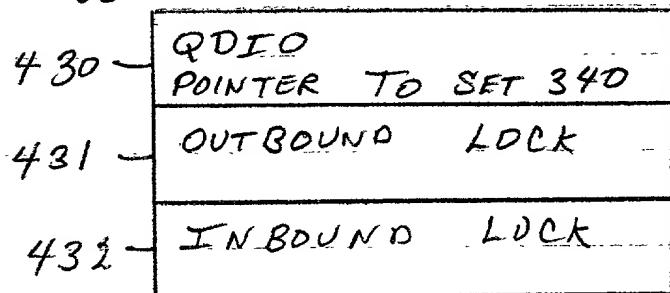


FIG. 5

QUEUE SET 340

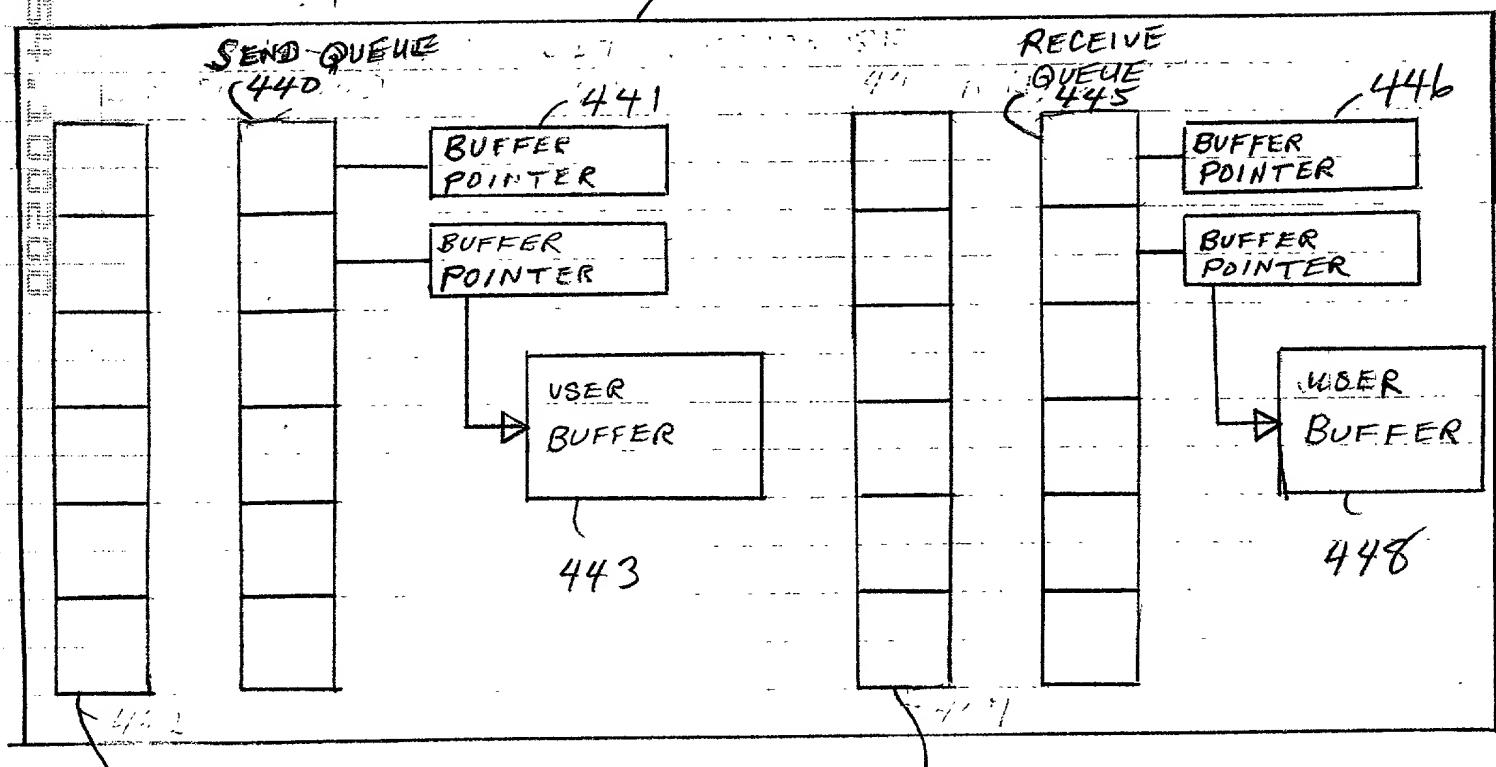


FIG. 6

STORAGE
LIST
STATUS
BLOCK
442

STORAGE
LIST
STATUS
BLOCK
447

243 USER BUFFER 244 TARGET IP ADDRESS

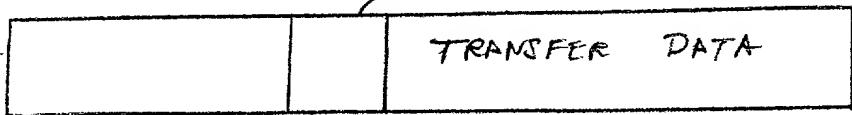


FIG. 7

312 SOURCE HASH DUPLICATE LIST ENTRIES

410-	LPAR=IP . SUBCHANNEL#
411-	STATUS
412-	POINTER TO QUEUE CONTROL 330
413-	NEXT POINTER TO 312
414-	PREVIOUS POINTER TO 311 OR 312

FIG. 8

322 TARGET HASH DUPLICATE LIST ENTRIES

420-	IP ADDRESS
421-	STATUS
422-	POINTER TO QUEUE CONTROL 330
423-	NEXT POINTER TO 322
424-	PREVIOUS POINTER TO 321 OR 322

FIG. 9

Docket No.
POU920000173US1

Declaration and Power of Attorney For Patent Application

English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

COMMUNICATIONS BETWEEN PARTITIONS WITHIN A LOGICALLY PARTITIONED COMPUTER

the specification of which

(check one)

is attached hereto.

was filed on _____ as United States Application No. or PCT International Application Number _____
and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)

Priority Not Claimed

(Number)	(Country)	(Day/Month/Year Filed)	<input type="checkbox"/>
(Number)	(Country)	(Day/Month/Year Filed)	<input type="checkbox"/>
(Number)	(Country)	(Day/Month/Year Filed)	<input type="checkbox"/>

I hereby claim the benefit under 35 U.S.C. Section 119(e) of any United States provisional

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

I hereby claim the benefit under 35 U. S. C. Section 120 of any United States application(s), or Section 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. Section 112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, C. F. R., Section 1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

(Application Serial No.)

(Filing Date)

(Status)

(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)

(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)

(patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. *(list name and registration number)*

Floyd A. Gonzalez, Reg. 26,732
Lynn L. Augspurger, Reg. 24,227
Lawrence D. Cutter, Reg. 28,501
William A. Kinnaman, Jr., Reg. 27,650
William B. Porter, Reg. 33,135
Andrew J. Wojnicki, Jr., Reg. 43,995
Christopher A. Hughes, Reg. 26,914
Edward A. Pennington, Reg. 32,588
John E. Hoel, Reg. 26,279
Joseph D. Redmond, Jr. Reg. 18,753

Send Correspondence to: **Floyd A. Gonzalez**
IBM Corporation
2455 South Road, P386
Poughkeepsie, New York 12601

Direct Telephone Calls to: *(name and telephone number)*
Floyd A. Gonzalez, (845) 433-1163

Full name of sole or first inventor Donald W. Schmidt	Date
Sole or first inventor's signature	
Residence 77 Vincent Lane, Stone Ridge, New York 12481	
Citizenship USA	
Post Office Address 77 Vincent Lane, Stone Ridge, New York 12481	

Full name of second inventor, if any John A. Aiken, Jr.	Date
Second inventor's signature	
Residence 10609 Chinkapin Place, Raleigh, North Carolina	
Citizenship USA	
Post Office Address 10609 Chinkapin Place, Raleigh, North Carolina 27613-6305	

Full name of third inventor, if any Frank W. Brice, Jr.	Date
Third inventor's signature	
Residence 152 Woodland Drive, Hurley, New York	
Citizenship USA	
Post Office Address 152 Woodland Drive, Hurley, New York 12443	

Full name of fourth inventor, if any Janet R. Easton	Date
Fourth inventor's signature	
Residence 78 Glasco Turnpike, Woodstock, New York	
Citizenship USA	
Post Office Address P.O. Box 56, Woodstock, New York 12498	

Full name of fifth inventor, if any Wolfgang Eckert	Date
Fifth inventor's signature	
Residence Brombergweg 19, 71155 Altdorf, Germany	
Citizenship German	
Post Office Address Brombergweg 19, 71155 Altdorf, Germany	

Full name of sixth inventor, if any Marcus Eder	Date
Sixth inventor's signature	
Residence Vikarweg 8B, 70567 Stuttgart, Germany	
Citizenship German	
Post Office Address Vikarweg 8B, 70567 Stuttgart, Germany	

Full name of seventh inventor, if any Steven G. Glassen	
Seventh inventor's signature	Date
Residence 11 Sherwood Drive, Wallkill, New York	
Citizenship USA	
Post Office Address 11 Sherwood Drive, Wallkill, New York 12589	

Full name of eighth inventor, if any Jeffrey P. Kubala	
Eighth inventor's signature	Date
Residence 10 Morgan Lane, Poughquag, New York	
Citizenship USA	
Post Office Address 20 Morgan Lane, Poughquag, New York 12570	

Full name of ninth inventor, if any Jeffrey M. Nick	
Ninth inventor's signature	Date
Residence 1957 Route 9W, West Park, New York	
Citizenship USAS	
Post Office Address P.O. Box 277, West Park, New York 12493	

Full name of tenth inventor, if any Jerry W. Stevens	
Tenth inventor's signature	Date
Residence 2101 Brixham Court, Raleigh, North Carolina	
Citizenship USA	
Post Office Address 2101 Brixham Court, Raleigh, North Carolina 27615	

Full name of eleventh inventor, if any Ambrose A. Verdibello, Jr.	
Eleventh inventor's signature	Date
Residence 64 Stanford Road, Millbrook, New York	
Citizenship USA	
Post Office Address 64 Stanford Road, Millbrook, New York 12545	

Full name of twelfth inventor, if any Harry M. Yudenfriend	
Twelfth inventor's signature	Date
Residence 1 Nob Hill Road, Poughkeepsie, New York	
Citizenship USA	
Post Office Address 1 Nob Hill Road, Poughkeepsie, New York 12603	

Full name of thirteenth inventor, if any Heinrich K. Lindner	
Thirteenth inventor's signature	Date
Residence Lilienstr. 14, 71088 Holzgerlingen, Germany	
Citizenship German	
Post Office Address Lilienstr. 14, 71088 Holzgerlingen, Germany	

Full name of fourteenth inventor, if any	
Fourteenth inventor's signature	Date
Residence	
Citizenship	
Post Office Address	